

Amendments to the Drawings:

The attached four formal drawing sheets replace the original sheets of formal drawings.

Amend the original formal drawings as follows:

Figure 1: Add an arrow pointing from the processor sets to the “failures” grid.

Place a block around all elements of the figure.

Figure 2: Add an arrow pointing from the processor sets to the “failures” grid.

Place a block around all elements of the figure.

Figure 3a: Add an arrow pointing from the processor sets to the “failures” grid.

Place a block around all elements of the figure.

Figure 3b: Add an arrow pointing from the processor sets to the “failures” grid.

Place a block around all elements of the figure.

Attachment: Replacement Drawing Sheets (4 sheets)

REMARKS

Claims 1-32 are pending. Claims 1 and 16 were amended to more particularly point out and distinctly claim the present invention. Claims 2, 8 and 17 were amended solely to address the 35 U.S.C. § 112, second paragraph, rejection. Claims 25-32 were added to further define the present invention. No new matter was added, as explained below.

Claims 1 and 16: The new limitation wherein the nodes “cooperate” with each other is disclosed on at least page 23, lines 20 and 23; page 24, line 10; page 25, line 2; and page 28, line 17. The new limitation of “means to allow one or more of the nodes to take over processing capacity of a node that becomes lost” is disclosed on at least page 24, lines 15-17; page 25, lines 1-4; page 27, line 24 through page 28, line 2; page 28, lines 28-30; and page 50, lines 2-3. The new limitation “the availability of the split processing system being greater than the availability of an unsplit system wherein all of the processors are located at a single node” is disclosed throughout the present application (see, for example, page 24, lines 9-13) and is an inherent limitation of the present invention which is directed to a split processing system for providing increased system availability.

Claims 25 and 26: The new limitation wherein “each node operates independently of the other nodes” is clearly disclosed on at least page 25, lines 12-14 of Applicants’ original specification.

Claims 27 and 28: The new limitation wherein “each node has less processors than the number of processors in the unsplit system” is clearly disclosed in at least Figs. 5a, 5b, 5c, 5e, and the corresponding disclosure in the specification.

Claims 29 and 31: The new limitation wherein “the split system has a specific number of failure modes which is less than the number of failure modes in an unsplit system wherein all of the processors are located at a single node” is clearly disclosed in at least page 23, lines 19-26.

Claims 30 and 32: The new limitation wherein “the number of failure modes is calculated from the number of nodes and spares of critical components” is clearly disclosed in at least Table 2 on page 18 and the text on page 4, lines 1-9.

For at least the reasons set forth below, withdrawal of all outstanding objections and rejections is respectfully requested.

Drawings

Figures 1, 2, 3a and 3b are being resubmitted in revised form. Each of these figures include a "failures" grid that relates to the respective figure. To better illustrate that the grid relates to the figures, an arrow was added pointing from each set of processors to the respective "failure" grid, and a block was added surrounding all of the figure elements.

35 U.S.C. § 112, second paragraph, rejection

Claims 2, 8 and 17 were amended to clarify the meaning of "at least a portion of a database." The amended phrase is believed to be clear. In view of the amendment, withdrawal of this rejection is respectfully requested.

Prior Art Rejection

Claims 1, 13-15, 16 and 22-24 were rejected under 35 U.S.C. § 103(a) as allegedly being unpatentable over Leymann et al. (hereafter, "Leymann"). Withdrawal of this rejection as it pertains to the amended claims is respectfully traversed.

1. Leymann

Leymann discloses a scheme for increasing interoperability of a workflow management system (WFMS) and a database management system (DBMS) (column 3, lines 25-26). The scheme involves splitting a database into a set of nodes with corresponding partitioned tables (Fig. 3; column 10, lines 63-67; column 11, lines 56-58; column 12, lines 2-9). Each node has a processor, and the node processors communicate with each other via a network (column 11, lines 23-25). Column 3, lines 47-58, which reads as follows, describes some of the ways that partitioned tables (i.e., a splitting process) can improve system performance:

Partitioned tables allow a program to work on part of the data at a time, while allowing concurrent access to other programs on other partitions. It becomes possible to put more frequently accessed data on faster devices.

More frequently accessed data can be separated from the remainder and can be put in a partition of its own and can use different device type. A single query to a partitioned database can initiate multiple parallel operations. These smaller queries run simultaneously on multiple processors accessing data in parallel. This reduces the elapsed time for a query.

Fig. 2 shows how a single logical table is split up and stored as a set of partitioned tables at their respective nodes. Partitioning keys of the individual records are used to determine which record gets stored in which partition table (column 11, lines 36-46). Each node contains a unique portion of the original database. Accordingly, if one of the nodes becomes lost for any reason, the remaining nodes cannot make up for the processing capacity of the lost node because the other nodes do not have the portion of the database that is in the lost node.

Splitting a database into a plurality of nodes, each node containing unique portions, may increase performance (as discussed in Leymann), but it does not necessarily increase system availability compared to an unsplit database system. Nowhere does Leymann disclose or suggest that there is greater availability in Leymann's split system compared to the original unsplit system. In fact, it is likely that there is less availability in Leymann as a result of splitting the database in the manner that Leymann performed the split (i.e., providing only unique portions of the database at each node). Leymann therefore provides no teaching of the present invention which seeks to increase availability compared to an unsplit system. In fact, since it is likely that there is less availability in Leymann as a result of splitting the database, Leymann likely teaches against the present invention.

Leymann discusses "availability" in numerous text portions, repeated below for convenience (underlining added for emphasis):

Text portion 1 (column 3, lines 8-17)

To cope with this new challenges the interoperation of the WFMS and the underlying database management system (DBMS) have to be improved. Specific issues are the issues of concurrency, parallelism and availability.

Techniques, such as hotpooling and application server clustering, have been proposed to improve performance and availability. However, no approach has yet been proposed to improve via the interoperation of the WFMSs and the underlying data base management systems (DBMS).

Text portion 2 (column 3, lines 41-51)

The introduction of parallel database technology within workflow management systems allows the latter to cope with the increasing requirements and expectations. Storing portions of this system repository in parallel databases leads to significant improvements with respect to parallelism, concurrency and availability. Parallel databases allow to work on part of the data at a time, cutting the time required for the operation to a manageable size. Partitioned tables allow a program to work on part of the data at a time, while allowing concurrent access to other programs on other partitions.

Text portion 3 (column 3, line 64 through column 4, line 1)

Especially introduction of inter-partition-parallel-relational-database technology, as introduced according to a further embodiment of the current invention, offers the greatest advantages with respect to parallelism, concurrency and availability.

Text portion 4 (column 5, lines 5-9)

Exploiting inter-parallel database technology the parallelized parts of the system repository actually are stored in physically different partitions providing a maximum of parallelism, concurrency and availability.

Text portion 5 (column 9, lines 50-56)

As already apparent from this system structure the database management system, holding the system repository, may become te [sic] bottleneck in terms of concurrency, parallelism and availability as the system repository is the focus of all data traffic; this conflict is due to the distributed approach of that WFMS on one hand and the central approach of DBMS on the other hand.

Text portion 6 (column 15, lines 25-29)

Exploiting inter-parallel database technology the parallelized parts (404, 405) of the system repository actually are stored in physically different partitions providing a maximum of parallelism, concurrency and availability.

The first text portion in Leymann merely discusses that techniques other than those employed by Leymann have been used to address availability, and highlights that Leymann is directed instead to improving interoperability.

The second text portion in Leymann discusses how parallel databases improve availability, but does not state that partitioned tables (which is the invention in Leymann) accomplishes this goal.

The third text portion in Leymann discusses how Leymann's invention offers the greatest advantage with respect to parallelism, concurrency and availability, but does not state that availability is increased. This text portion merely states that given these three goals, Leymann's invention offers the greatest advantage in balancing all three goals.

The fourth text portion in Leymann discusses how Leymann's invention maximizes parallelism, concurrency and availability, but does not state that availability is increased. This text portion merely states that given these three goals, Leymann's invention offers a maximized solution. The fourth text portion is thus similar to the third text portion regarding the lack of a disclosure of increasing availability.

The fifth text portion has no disclosure of increasing availability. Availability is merely being mentioned in a general context.

The sixth text portion in Leymann is identical in content to the fourth text portion, and thus also lacks a disclosure of increasing availability.

Leymann thus has no disclosure of increasing availability. This lack of disclosure is not unexpected because the type of database splitting that occurs in Leymann likely reduces system availability, even though it may increase system performance.

In sum, Leymann merely discloses a database splitting process, and thus is not relevant to the present invention, or to the claim limitations directed to increased availability.

2. Patentability of claims 1 and 16 over Leymann

Claims 1 and 16 read as follows (underlining added for emphasis):

1. A split processing system comprising:

- (a) a plurality of nodes, each node including one or more processors
- (b) a communication network that allows the one or more processors at each of the nodes to cooperate with each other; and
- (c) means to allow one or more of the nodes to take over processing capacity of a node that becomes lost,
the availability of the split processing system being greater than the availability of an unsplit system wherein all of the processors are located at a single node.

16. A split processing system comprising:
- (a) a plurality of nodes, each node including:
 - (i) a processor subsystem including at least one processor,
 - and
 - (ii) an operating system;
 - (b) a communication network that allows the one or more processors at each of the nodes to cooperate with each other; and
 - (c) means to allow one or more of the nodes to take over processing capacity of a node that becomes lost,
the availability of the split processing system being greater than the availability of an unsplit system wherein all of the processors are located at a single node.

Regarding clause (c) in claims 1 and 16, as discussed above, Leymann cannot allow one of the nodes to take over processing capacity of a node that becomes lost because each node contains a unique portion of the database. Therefore, Leymann does not disclose or suggest clause (c).

Regarding the “availability” limitations in claims 1 and 16, as discussed above, nowhere does Leymann disclose or suggest that there is greater availability in Leymann’s split system compared to the original unsplit system. In fact, it is likely that there is less availability in Leymann as a result of splitting the database in the manner that Leymann performed the split (i.e., providing only unique portions of the database at each node). Leymann therefore provides no teaching of the increasing availability compared to an unsplit system. In fact, since it is likely that there is less availability in Leymann as a result of splitting the database, Leymann likely teaches against the increased availability limitations in claims 1 and 16.

For at least the reasons set forth above, withdrawal of the prior art rejection is respectfully requested.

3. Patentability of dependent claims 27-28

Leymann does not disclose or suggest a split processing system wherein each node has less processors than the number of processors in the unsplit system. Leymann is completely silent as to how many processors are present before and after the system splitting. Furthermore, these claims are believed to be allowable because they depend upon respective allowable independent claims.

4. Patentability of dependent claims 29 and 31

In the outstanding Office Action, the Examiner admits that Leymann does not explicitly teach the original claim limitations directed to reduced failure modes. However, the Examiner states that this limitation would have been obvious in view of Leymann's disclosure of improving performance. Stated simply, the Examiner presumes that as performance goes up, the number of failure modes (which is generally inversely related to availability) goes down. However, this is an incorrect assumption. Improving performance does not necessarily reduce failure modes. In fact, the manner in which Leymann improves performance, namely, by splitting a database into unique portions, creating a plurality of nodes, and placing a unique portion in each node, will likely increase the number of failure modes. Leymann likely teaches away from claims 29 and 31. Furthermore, these claims are believed to be allowable because they depend upon respective allowable independent claims.

5. Patentability of dependent claims 30 and 32

Leymann has no disclosure of how to calculate failure modes or availability. Also, Leymann does not even discuss "spares." Furthermore, these claims are believed to be allowable because they depend upon respective allowable independent claims.

6. Patentability of dependent claims 2-15, 17-24 and new dependent claims 25-26

The dependent claims are believed to be allowable because they depend upon respective allowable independent claims, and because they recite additional patentable steps.

Conclusion

Insofar as the Examiner's rejections were fully addressed, the instant application is in condition for allowance. Issuance of a Notice of Allowability of all pending claims is therefore earnestly solicited.

Application No. 10/700,410
Reply to Office Action of October 5, 2005

Respectfully submitted,

WILBUR H. HIGHLEYMAN et al.

February 28, 2006 By: Clark Jablon
(Date)

CLARK A. JABLON
Registration No. 35,039
AKIN GUMP STRAUSS HAUER & FELD LLP
One Commerce Square
2005 Market Street - Suite 2200
Philadelphia, PA 19103-7086
Telephone: (215) 965-1200
Direct Dial: (215) 965-1293
Facsimile: (215) 965-1210
E-Mail: cjablon@akingump.com

CAJ:gem

Enclosure: Drawing Replacement Sheets (4 sheets)